
A Markov Jump Process for More Efficient Hamiltonian Monte Carlo

Andrew B Berger
Redwood Center
UC Berkeley

Mayur Mudigonda
Redwood Center
UC Berkeley

Michael R DeWeese
Redwood Center
UC Berkeley

Jascha Sohl-Dickstein
Stanford University¹

Abstract

In most sampling algorithms, including Hamiltonian Monte Carlo, transition rates between states correspond to the probability of making a transition in a single time step, and are constrained to be less than or equal to 1. We derive a Hamiltonian Monte Carlo algorithm using a continuous time Markov jump process, and are thus able to escape this constraint. Transition rates in a Markov jump process need only be non-negative. We demonstrate that the new algorithm leads to improved mixing for several example problems, both by evaluating the spectral gap of the Markov operator, and by computing autocorrelation as a function of compute time. We release the algorithm as an open source Python package.

1 Introduction

Efficient sampling is a challenge in many tasks involving high dimensional probabilistic models, in a diversity of fields. For example, sampling is commonly required to train a probabilistic model, to evaluate the model’s performance, to perform inference, and to take expectations under the model [1].

In this paper we introduce a method for more efficient sampling, by making Markov transitions in continuous rather than discrete time. This allows transitions into lower probability states to occur more often, with a shorter time spent for each visit, and thus allows for more rapid exploration of state space. We apply this approach to develop a novel Hamiltonian Monte Carlo (HMC) sampling algorithm. Finally, we demonstrate the effectiveness of this approach by comparing both spectral gaps and autocorrelation on several example problems.

1.1 Discrete time sampling

Most sampling algorithms involve transitioning between states in discrete time steps, with a fixed interval. In this discrete-time framework, the transition rates out of a state must sum to 1 and be non-negative. In fact, the popular Metropolis-Hastings acceptance rule [2] for Markov Chain Monte Carlo (MCMC) works well because it maximizes the transition rate between a pair of states, subject to this constraint and to detailed balance. As we will see, however, this constraint on transition rates limits performance, and better mixing can be achieved by allowing transition rates larger than 1.

1.2 Markov jump process

A Markov process can also be expressed in continuous time, in which case the only restriction on the transition rates between distinct states is that they be non-negative. In continuous time, the rate of transition from a state j into a state $i \neq j$ is given by Γ_{ij} , and the rate of change of the probability p_i of state i is

$$\frac{\partial p_i}{\partial t} = \sum_j \Gamma_{ij} p_j, \quad (1)$$

where $\Gamma_{ij} \geq 0$ for $\forall i \neq j$, and we use the convention $\Gamma_{jj} = -\sum_{i \neq j} \Gamma_{ij}$.

If a particle evolves in a system of this form it makes stochastic transitions between a set of discrete states in continuous time. Each transition is governed by a Poisson process. Neglecting other states, the waiting time w_{ij} for a transition from state j to state i will be drawn from an exponential distribution, $P(w_{ij}) = \Gamma_{ij} \exp(-\Gamma_{ij} w_{ij})$.

To simulate the system, waiting times w_{kj} are generated for all candidate states k , and the shortest waiting time $i = \operatorname{argmin}_k w_{kj}$ is chosen. We call this shortest waiting time the holding time. A transition is then performed to state i after a delay of length w_{ij} .

¹Currently at Google Research.

A system that evolves in this way is known as a Markov jump process [3]. Markov jump processes have been used to model physical systems, such as chemical reactions, which are well described by stochastic transitions between discrete states [4]. Work has also been done on efficient sampling of trajectories in Markov jump processes [5] and the statistical properties of these trajectories [6].

To our knowledge, Markov jump processes have not previously been applied to general purpose Monte Carlo sampling, though see [7] where a Markov jump process is used to sample from a posterior distribution over model graph structure. One barrier to using a Markov jump process for general Monte Carlo is that it is necessary to compute transition rates to all possible target states i from the current state j . As we will see, however, for HMC we only need to consider a small number of target states. **The primary contribution of this paper is to use a Markov jump process to develop a faster HMC algorithm.**

1.2.1 Relationship to importance sampling

As will be elaborated in Section 2.4, samples from a Markov jump process can be generated by using an equivalent discrete time process $\hat{\mathbf{T}}$ to generate the same distribution over state *sequences*, and then resampling according to each state's holding time. From this perspective, the process of sampling from a Markov jump process can be seen as a realization of importance sampling, with a particularly unusual proposal distribution. The equivalent discrete time process $\hat{\mathbf{T}}$ defines the importance sampling proposal distribution, and the holding times provide the importance weights.

The discrete time distribution $\hat{\mathbf{p}}$ generated by $\hat{\mathbf{T}}$ will tend to be more similar to a uniform distribution than \mathbf{p} , and the corresponding Markov chain $\hat{\mathbf{T}}$ will thus tend to mix more quickly than a typical discrete time sampler. For instance, $\hat{\mathbf{T}}$ has no self-transitions, so unlike in a standard Metropolis-Hastings algorithm there is no sample rejection, and as a result there is likely to be less wasted computation. Qualitatively, the equivalent discrete time process $\hat{\mathbf{T}}$ can be expected to visit low probability states far more frequently than an unweighted sampler. Those states will just have very short holding times, and be assigned very small importance weights. This will allow it to more rapidly explore the state space.

1.3 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) [8, 9] is the state-of-the-art, general purpose Monte Carlo algorithm for sampling from a distribution $\pi(\mathbf{x})$ over a continuous

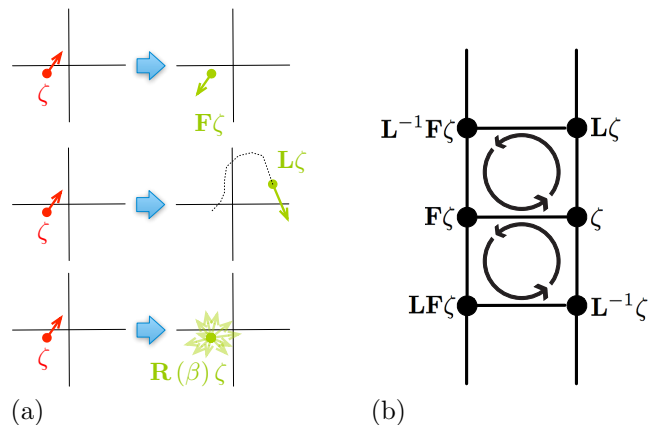


Figure 1: (a) The action of operators involved in Hamiltonian Monte Carlo (HMC). The base of each red or green arrow represents the position \mathbf{x} , and the length and direction of each of these arrows represents the momentum \mathbf{v} . The flip operator \mathbf{F} reverses the momentum. The leapfrog operator \mathbf{L} approximately integrates Hamiltonian dynamics. The trajectory taken by \mathbf{L} is indicated by the dotted line. The randomization operator \mathbf{R} replaces the momentum. (b) The ladder of discrete states generated by the leapfrog (\mathbf{L}) and flip (\mathbf{F}) operators. Application of \mathbf{F} corresponds to movement across the ‘rungs’ of the ladder. Application of \mathbf{L} corresponds to movement up the right side of the ladder, or down the left. Inset arrows illustrate closed loops of constant total probability flow under our chosen rate (Equation 3)

state space $\mathbf{x} \in \mathbb{R}^N$.

HMC utilizes the same Hamiltonian dynamics that govern the evolution of a physical system – for instance a marble rolling in a swimming pool – to rapidly traverse long distances in state space. In HMC the state space is first extended to include auxiliary momentum variables $\mathbf{v} \in \mathbb{R}^N$ with distribution $\pi(\mathbf{v})$, such that the joint state space over position and momentum is $\zeta = \{\mathbf{x}, \mathbf{v}\}$, with joint distribution $\pi(\zeta) = \pi(\mathbf{x})\pi(\mathbf{v})$. An analogy is then made between \mathbf{x} and physical position (e.g. the position of the marble), between \mathbf{v} and physical momentum (the momentum of the marble), and between $(-\log \pi(\mathbf{x}))$ and potential energy (the height of the swimming pool at position \mathbf{x}). Since physical dynamics conserve energy, they can generate very long trajectories in state space while remaining on a constant probability contour of $\pi(\mathbf{zeta})$.

HMC is thus able to move very long distances in state space in a single update step.

1.3.1 Discrete state space Ladder

As introduced in [10] and illustrated in Figure 1, HMC can be viewed in terms of transitions on a discrete state space ladder. This state ladder is formulated by expressing the action of HMC on a sampling particle in terms of three operators. The leapfrog integration operator, \mathbf{L} , approximately integrates Hamiltonian dynamics for a fixed number of time steps and a fixed step length.

The momentum flip operator, \mathbf{F} , reverses a particle's direction of travel along a contour by flipping its momentum. The momentum randomization operator \mathbf{R} redraws the momentum vector from $\pi(\mathbf{v})$, and moves a particle onto a new state space ladder.

This perspective suggests a powerful formalism, effectively discretizing the state space, and it illuminates the structure² of HMC. Throughout this work, we refer to the structure generated by the operators as a state ladder. As illustrated in Figure 1, \mathbf{L} causes movement up the right side of a state ladder and down the left side, whereas \mathbf{F} causes horizontal movement across the rungs of a ladder. A trajectory can be exactly reversed by reversing the momentum, integrating Hamiltonian dynamics, and reversing the momentum again. As can be seen in Figure 1, this corresponds to making a loop on the state space ladder, and it implies that $\mathbf{FLFL} = \mathbf{I}$, where \mathbf{I} is the identity operator. \mathbf{R} causes movement off of the current state ladder and onto a new one. Both \mathbf{F} and \mathbf{L} are volume preserving, which will eliminate the need to consider volume changes when computing Markov transition rates.

If we only allow transitions between states that are connected on the state ladder (Figure 1), then transitions can only occur between ζ and three other states ($\mathbf{L}^{-1}\zeta$, $\mathbf{F}\zeta$, $\mathbf{L}\zeta$). This makes HMC well matched to Markov jump processes, since only a small number of transitions need be considered.

1.3.2 Current research

The development of improved methods for performing HMC is an important area of active research. These include the use of shadow Hamiltonians that are more closely conserved by the approximate Hamiltonian integrator [11], Riemann manifold HMC [12] and other investigation of its geometry [13], quasi-Newton HMC [14], Hilbert space HMC [15], parameter adapta-

²Between momentum randomizations, HMC acts in a manner isomorphic to the Dihedral group of, in general, infinite order. The HMC state ladder, and thus the order, is generally infinite because trajectories through state space produced by Hamiltonian dynamics are almost never closed.

tion techniques [16], Hamiltonian annealed importance sampling [17], split HMC [18], tempered trajectories [9], novel discrete time transition rules [10, 19, 20], stochastic gradient variants on HMC [21], HMC for approximate Bayesian computation [22], and new approximate Hamiltonian integrators [23].

2 Markov jump Hamiltonian Monte Carlo

We now discuss various aspects of our MJHMC sampler, schematized in Figure 2.

2.1 Continuous time transition rates on HMC state ladder

A Markov process must satisfy *two conditions* [24] to sample from a target distribution. The first is ergodicity, which requires that the process will eventually explore the full state space; this is typically straightforward to satisfy.

The second condition is that the target distribution must be a fixed point of the Markov process. This is usually achieved via dynamics that satisfy detailed balance matched to the state probabilities of the target distribution.

2.1.1 Closed loops preserve the fixed point distribution

Markov transition rates $\tilde{\Gamma}(\zeta' | \zeta)$ that preserve $\pi(\zeta)$ as a fixed point can be constructed from closed loops of constant probability flow in state space.

For closed loops to have constant probability flow, the flow $r(\zeta_i | \zeta_j) = \Gamma(\zeta_i | \zeta_j) \pi(\zeta_j)$ from state j to i must be identical for each link in the loop. This is analogous to Kirchoff's current law for electrical circuits – constant probability flow in all loops implies that the net probability flow into any state equals the net flow out of that state.

2.1.2 Choosing transition rates

For the case of HMC, we set the loops to be between the “rungs” of the state ladder, as illustrated in Figure 1. The loop balance condition for each closed loop becomes

$$\tilde{r}(\mathbf{F}\zeta | \zeta) = \tilde{r}(\mathbf{L}\mathbf{F}\zeta | \mathbf{F}\zeta) = \tilde{r}(\mathbf{L}^{-1}\zeta | \mathbf{L}\mathbf{F}\zeta) = \tilde{r}(\zeta | \mathbf{L}^{-1}\zeta), \quad (2)$$

$$\begin{aligned} \pi(\zeta) \tilde{\Gamma}(\mathbf{F}\zeta | \zeta) &= \pi(\mathbf{F}\zeta) \tilde{\Gamma}(\mathbf{L}\mathbf{F}\zeta | \mathbf{F}\zeta) \\ &= \pi(\mathbf{L}\mathbf{F}\zeta) \tilde{\Gamma}(\mathbf{L}^{-1}\zeta | \mathbf{L}\mathbf{F}\zeta) \\ &= \pi(\mathbf{L}^{-1}\zeta) \tilde{\Gamma}(\zeta | \mathbf{L}^{-1}\zeta). \end{aligned} \quad (3)$$

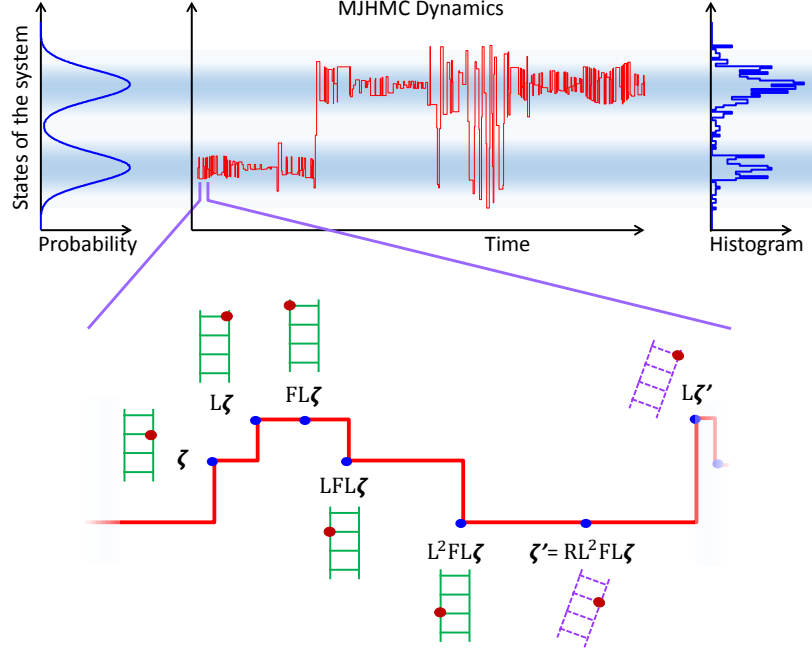


Figure 2: Illustration of Markov Jump HMC sampling dynamics. The red curve represents a particle trajectory for 400 time steps. Blue shading indicates the probability density, plotted at the left, and an empirical histogram of the samples is shown at the right. The inset blowup at the bottom shows how movement of the sampling particle corresponds to transitions on the state ladder, using the symbolic and graphical conventions described in Figure 1. Note that the sampling particle dwells in a position for a duration related to the probability density of that state relative to its neighbors. We have indicated the transition from one state ladder (vertical green ladder) to a new ladder (angled purple ladder) following a momentum randomization event, resulting in a new state labeled ζ' .

In order to satisfy this condition, we set the transition rates to be

$$\tilde{\Gamma}(\zeta' | \zeta) = \begin{cases} \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} & \zeta' = \mathbf{L}\zeta \\ \left[\frac{\pi(\mathbf{L}\mathbf{F}\zeta)}{\pi(\mathbf{F}\zeta)} \right]^{\frac{1}{2}} & \zeta' = \mathbf{F}\zeta \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

One can verify by direct substitution that these transition rates satisfy Equation 3. The transition rates for the full ladder consist of a sum over the transition rates for each loop.

2.1.3 Opposing flows cancel

As can be seen in Figure 1, adjacent loops make flip transitions across the “rungs” of the ladder in opposite directions. After summing over all loops, the net transitions across the ladder approximately cancel.

This allows us to reduce the flip rates in both directions, such that the flip rate in one direction is zero. The final rate of flip transitions in our algorithm will thus be

$$\Gamma(\mathbf{F}\zeta | \zeta) = \tilde{\Gamma}(\mathbf{F}\zeta | \zeta) - \min \left[\tilde{\Gamma}(\mathbf{F}\zeta | \zeta), \tilde{\Gamma}(\zeta | \mathbf{F}\zeta) \right], \quad (5)$$

$$= \max \left[0, \tilde{\Gamma}(\mathbf{F}\zeta | \zeta) - \tilde{\Gamma}(\zeta | \mathbf{F}\zeta) \right], \quad (6)$$

$$= \max \left[0, \left[\frac{\pi(\mathbf{L}\mathbf{F}\zeta)}{\pi(\mathbf{F}\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}\mathbf{F}\mathbf{F}\zeta)}{\pi(\mathbf{F}\mathbf{F}\zeta)} \right]^{\frac{1}{2}} \right], \quad (7)$$

$$= \max \left[0, \left[\frac{\pi(\mathbf{L}^{-1}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \right], \quad (8)$$

where the final step relies on the observations that $\mathbf{F}\mathbf{F}\zeta = \zeta$, and that $\pi(\zeta) = \pi(\mathbf{F}\zeta)$ [10]. In practice, $\pi(\mathbf{L}^{-1}\zeta)$ will typically already be available (up to a shared normalization constant) from the preceding Markov transition, and will not need to be computed.

Due to discretization error, the leapfrog integrator for Hamiltonian dynamics only approximately conserves

probability. Equation 8 shows that the residual flow across the ladder stems from this discretization error of the leapfrog integrator. This is completely analogous to the cause of momentum flips in standard HMC.

2.2 Momentum randomization

In discrete time HMC the momentum is periodically corrupted with noise. If this was not done, then sampling would be restricted to a single state ladder, and mixing between state ladders would not occur. In order to accomplish the same end in continuous time, we jump to a state $\mathbf{R}\zeta$ with a constant transition rate β . A transition to $\mathbf{R}\zeta$ corresponds to replacing the momentum \mathbf{v} with a new draw from $\pi(\mathbf{v})$. The transition rate from \mathbf{v} to a particular \mathbf{v}' is thus $\beta\pi(\mathbf{v}')$. It can be seen by substitution that this rate satisfies detailed balance.

2.3 Final transition rates

Combining the transition rates derived in Sections 2.1.2, 2.1.3, and 2.2,

$$\Gamma(\zeta' | \zeta) = \begin{cases} \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} & \zeta' = \mathbf{L}\zeta \\ \max \left[0, \left[\frac{\pi(\mathbf{L}^{-1}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \right] & \zeta' = \mathbf{F}\zeta \\ \beta & \zeta' = \mathbf{R}\zeta \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

We verify that these transition rates satisfy the balance condition for $\pi(\zeta)$ in Appendix A. The third line is a slight abuse of notation since $\mathbf{R}\zeta$ does not correspond to a single fixed state, but rather indicates that the momentum is replaced by a new draw from $\pi(\mathbf{v})$, where this replacement is triggered by a Poisson process with rate β . Note that as in [10] these dynamics do not satisfy detailed balance, and can be expected to mix more quickly as a result [25].

2.4 System time vs compute time

We have described continuous time dynamics in terms of a system, or simulation, time. However, when applying this sampler to a real problem it is its performance as measured relative to compute time that matters. Here we show how to relate the continuous time dynamics of the Markov jump process to a discrete time Markov process, with an approximately fixed computational cost per time-step.

First we observe that there is a discrete time Markov process describing only the sequence of visited states, thus neglecting the holding time spent in each state.

For notational convenience we represent Markov processes using matrix notation in this section. The update rule for this Markov process can be written

$$\hat{\mathbf{p}}^{\tau+1} = \hat{\mathbf{T}}\hat{\mathbf{p}}^{\tau}, \quad (10)$$

$$\hat{T}_{ij} = \begin{cases} \prod_{k \neq j} \frac{\Gamma_{ij}}{\Gamma_{ij} + \Gamma_{ik}} & i \neq j \\ 0 & i = j \end{cases} \quad (11)$$

where the matrix $\hat{\mathbf{T}}$ is the Markov transition kernel, and the vector $\hat{\mathbf{p}}^{\tau}$ is the probability distribution over system states at timestep τ . The computational cost of each time step is roughly constant under this Markov chain, since each step requires computing the transition rate to all possible next states in Equation 4.

The current and fixed point distributions $\hat{\mathbf{p}}$ and $\hat{\pi}$ under this process can be related to the corresponding distributions \mathbf{p} and π under the Markov jump process by scaling by the expected holding time,

$$\pi = \frac{1}{Z} \mathbf{D} \hat{\pi}, \quad (12)$$

$$\mathbf{p} = \frac{1}{Z} \mathbf{D} \hat{\mathbf{p}}, \quad (13)$$

where \mathbf{D} is a diagonal matrix with the expected holding times for each state on the diagonal, $D_{jj} = \frac{1}{\sum_{i \neq j} \Gamma_{ij}}$, and Z is a normalization constant [26]. Similarly, the evolution of \mathbf{p} relative to these discrete time steps can be expressed by scaling $\hat{\mathbf{p}}$ in Equation 10 by the holding time,

$$\mathbf{p}^{\tau+1} = \mathbf{D} \hat{\mathbf{T}} \mathbf{D}^{-1} \mathbf{p}^{\tau}, \quad (14)$$

$$\mathbf{p}^{\tau+1} = \mathbf{T} \mathbf{p}^{\tau}, \quad (15)$$

where $\mathbf{T} = \mathbf{D} \hat{\mathbf{T}} \mathbf{D}^{-1}$ describes the discrete time evolution of the samples. Since \mathbf{T} and $\hat{\mathbf{T}}$ are related to each other by a similarity transform, they share identical eigenvalues. In order to evaluate the spectral gap, and thus the mixing time, of the Markov jump process in terms of computational time, it is thus sufficient to compute the spectral gap of $\hat{\mathbf{T}}$. We do this for randomly generated toy systems in Section 3.1 and Figure 3 and show that MJHMC has superior spectral gap characteristics, indicating more efficient mixing.

2.5 Algorithm

Here, we briefly summarize the Markov Jump HMC algorithm for generating N samples in pseudocode. As with all HMC sampling algorithms, an energy function $E(x)$ (equivalent to $-\log \pi(\mathbf{x})$ plus a constant) and its gradient are required. The three hyperparameters are the leapfrog step size ϵ , the number of leapfrog steps per sampling step M ; and the momentum corruption rate β .

Note that computation of $\mathbf{L}^{-1}\zeta$ is only necessary when the last transition made was a momentum flip or randomization. The number of times the gradient is evaluated in an MJHMC sampling step is comparable to that of standard HMC.

Algorithm 1: Markov Jump Process HMC

```

input :  $\epsilon, M, \beta, E(x), \nabla E(x), N$ 
output:  $N$  samples
1  $\zeta_0 \leftarrow$  Randomly initialized ;
2 for  $i \leftarrow 1$  to  $N$  do
3   Calculate states  $\mathbf{L}\zeta_{i-1}, \mathbf{F}\zeta_{i-1}, \mathbf{L}^{-1}\zeta_{i-1}$  ;
4   Compute  $E(\zeta_{i-1}), E(\mathbf{L}\zeta_{i-1}), E(\mathbf{F}\zeta_{i-1}),$ 
      $E(\mathbf{L}^{-1}\zeta_{i-1})$  ;
5   Compute transition rates  $\Gamma_L, \Gamma_F, \Gamma_R$  using
     Equation (9) ;
6   Draw waiting times  $w_L, w_F, w_R$  from an
     exponential distribution, using rates of  $\Gamma_L,$ 
      $\Gamma_F, \Gamma_R$  respectively;
7   Record holding time for  $\zeta_{i-1}$ .
      $h_{i-1} \leftarrow \min(w_L, w_F, w_R)$  ;
8   Set  $\zeta_i$  to whichever of  $\mathbf{L}\zeta, \mathbf{F}\zeta, \mathbf{R}\zeta$  had the
     shortest waiting time ;
9 end
10 Resample all  $\zeta_i$  using holding times  $h_i$  as
    importance weights ;
    
```

3 Experimental results

3.1 Spectral gap on HMC state ladder

The convergence rate of a Markov process to its steady state is given by its spectral gap [27]. This is the difference in the magnitude of the two largest eigenvalues. We numerically compute this value for randomly generated toy problems in order to compare our mixing rate to that of standard HMC. As all HMC algorithms randomize momentum in nearly the same way, it is expected that their mixing time over a single state ladder is representative of their mixing time over the entire state space. To achieve analytic tractability we restrict our attention to finite state ladders. To avoid edge effects, we attach the top and bottom rungs of the ladder to each other, so that the ladder forms a loop and $\mathbf{L}^k\zeta = \zeta$, where k is the number of distinct rungs. We evaluate the eigenvalues on each state ladder using the similarity relationship to a discrete time Markov chain in Equation 14.

A comparison of such spectral gaps between Markov Jump HMC and standard HMC is illustrated in Figure 3 as a function of state ladder size. We draw the energy for each ‘rung’ of the state space ladder from a unit norm Gaussian distribution, and average across

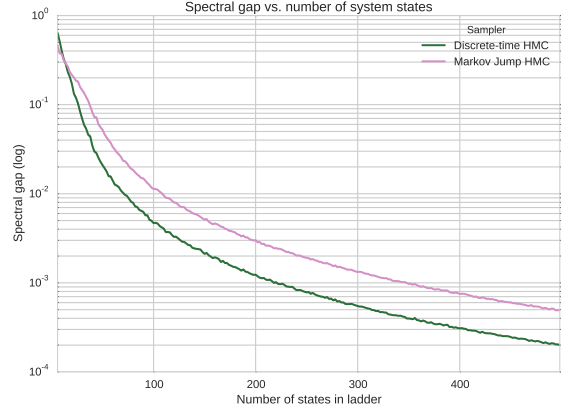


Figure 3: Comparison of mixing performance of Markov jump HMC (MJHMC) and standard discrete time HMC. Spectral gap versus size of state ladder. For large state ladder sizes, MJHMC is better by half an order of magnitude.

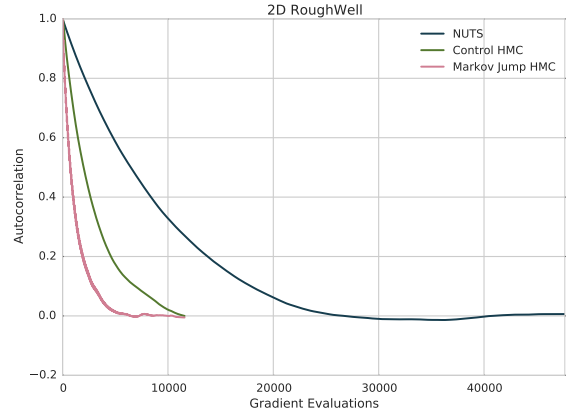


Figure 4: Autocorrelation versus number of gradient evaluations for standard HMC and MJHMC for the Rough Well distribution. The hyperparameters found by Spearmin for MJHMC are $\epsilon = 3.0, \beta = 0.012314, M = 25$ and for control HMC $\epsilon = 0.591686, \beta = 0.429956, M = 25$.

250 draws for each ladder size. Figure 3 thus shows performance averaged over many randomly generated energy landscapes. MJHMC mixes faster (has a larger spectral gap) for all except the smallest state space sizes.

3.2 Autocorrelation on rough well distribution

Explicit computation of mixing time for most problems is computationally intractable. It is common to instead use the rate at which the sample autocorrelation

approaches zero as a proxy. As illustrated in Figure 4, we compare autocorrelation traces for MJHMC with standard HMC and NUTS on the rough well distribution, and find that MJHMC performs significantly better.

Our results, illustrated in Figure 4, indicate that MJHMC significantly outperforms standard HMC.

3.2.1 Energy function

The chosen energy function was the ‘rough well’ distribution from [10]. This distribution provides a good test case because it is as simple as possible, while also presenting both well understood and significant challenges to HMC-style samplers. Its energy function is

$$E(x) = \frac{1}{2\sigma_1^2}(x_1^2 + x_2^2) + \cos\left(\frac{\pi x_1}{\sigma_2}\right) + \cos\left(\frac{\pi x_2}{\sigma_2}\right), \quad (16)$$

where $\sigma_1 = 100$ and $\sigma_2 = 4$. Although this distribution is well conditioned everywhere, the sinusoids cause it to have a ‘rough’ surface, such that it requires many leapfrog steps to traverse the quadratic well while maintaining a reasonable discretization error.

3.2.2 Hyperparameter selection

As HMC samplers are sensitive to the choice of hyperparameters, we searched for optimal settings of the hyperparameters for all samplers with the Spearmint package [28].

For each hyperparameter setting, we computed the autocorrelation as a function of number of gradient evaluations n , $C(n)$. We then fit a function $\rho(n; r)$ to this in terms of a complex coefficient $r \in \mathbb{C}$,

$$\hat{r} = \underset{r}{\operatorname{argmin}} \|\rho(n; r) - C(n)\|^2 \quad (17)$$

$$\rho(n; r) = \operatorname{Re}[\exp(rn)] \quad (18)$$

where the imaginary portion of r corresponds to an oscillatory rate, and the real part corresponds to the decay rate towards 0 autocorrelation. We use $\operatorname{Re}(r)$ as the Spearmint objective.

We provide additional figures to provide support for our hyperparameter search results in Appendix Figures A.1, A.2, and A.4.

NUTS self-tunes its hyperparameters during burn-in, so we did not perform a hyperparameter search for NUTS.

4 Discussion

We introduced an algorithm, Markov Jump Hamiltonian Monte Carlo (MJHMC), in which the state transitions in Hamiltonian Monte Carlo sampling occur as Poisson processes in continuous time, rather than at discrete time steps. We demonstrated that this algorithm led to improved mixing performance, as measured by explicit computation of the spectral gap, by the autocorrelation of the sampler on a simple but challenging distribution.

5 Acknowledgements

MRD and MM were both supported by NSF grant IIS-1219199 to Michael R. DeWeese. MM was also partially supported by NGA grant HM1582-081-0007 to Bruno Olshausen and NSF grant IIS-1111765 to Bruno Olshausen. ABB was partially supported by a Berkeley Physics Undergraduate Research Scholars Program Fellowship. This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under contract number W911NF-13-1-0390.

References

- [1] DJC MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, 2003.
- [2] W Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, January 1970.
- [3] Erhan Cinlar. *Introduction to stochastic processes*. Courier Corporation, 2013.
- [4] David F Anderson and Thomas G Kurtz. Continuous time Markov chain models for chemical reaction networks. In *Design and analysis of biomolecular circuits*, pages 3–42. Springer, 2011.
- [5] Vinayak Rao and Yee Whye Teh. Fast MCMC sampling for Markov jump processes and continuous time Bayesian networks. *arXiv preprint arXiv:1202.3760*, 2012.
- [6] Philipp Metzner, Christof Schütte, and Eric Vanden-Eijnden. Transition path theory for Markov jump processes. *Multiscale Modeling & Simulation*, 7(3):1192–1219, 2009.
- [7] U Grenander and MI Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1994.
- [8] S Duane, AD Kennedy, BJ Pendleton, and D Roweth. Hybrid monte carlo. *Physics letters B*, 1987.
- [9] Radford M Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, January 2010.

- [10] Jascha Sohl-Dickstein, Mayur Mudigonda, and Michael DeWeese. Hamiltonian Monte Carlo without detailed balance. In *Proceedings of the 31st International Conference on Machine Learning*, pages 719–726, 2014.
- [11] JA Izaguirre and SS Hampton. Shadow hybrid Monte Carlo: an efficient propagator in phase space of macromolecules. *Journal of Computational Physics*, 2004.
- [12] Mark Girolami and Ben Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, March 2011.
- [13] M J Betancourt, Simon Byrne, Samuel Livingstone, and Mark Girolami. The Geometric Foundations of Hamiltonian Monte Carlo. *arXiv preprint arXiv:1410.5110*, 2014.
- [14] Y Zhang and C Sutton. Quasi-Newton Markov chain Monte Carlo. 2011.
- [15] A Beskos, FJ Pinski, J. M. Sanz-Serna, and A. M. Stuart. Hybrid monte carlo on hilbert spaces. *Stochastic Processes and their Applications*, 2011.
- [16] Z Wang, S Mohamed, and D Nando. Adaptive Hamiltonian and Riemann Manifold Monte Carlo. *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, 2013.
- [17] Jascha Sohl-Dickstein and Benjamin J. Culpepper. Hamiltonian Annealed Importance Sampling for partition function estimation. *Redwood Center Technical Report*, (arXiv 1205.1925), May 2012.
- [18] B Shahbaba, S Lan, WO Johnson, and RM Neal. Split hamiltonian monte carlo. *Statistics and Computing*, 2011.
- [19] Jascha Sohl-Dickstein. Hamiltonian Monte Carlo with Reduced Momentum Flips. *Redwood Center Technical Report*, (arXiv 1205.1939), May 2012.
- [20] Cédric M Campos and J M Sanz-Serna. Extra Chance Generalized Hybrid Monte Carlo. *Journal of Computational Physics*, 281:365–374, 2015.
- [21] Tianqi Chen, Emily B Fox, and Carlos Guestrin. Stochastic gradient hamiltonian monte carlo. *arXiv preprint arXiv:1402.4102*, 2014.
- [22] Edward Meeds, Robert Leenders, and Max Welling. Hamiltonian ABC. *arXiv preprint arXiv:1503.01916*, 2015.
- [23] Wei-Lun Chao, Justin Solomon, Dominik L Michels, and Fei Sha. Exponential Integration for Hamiltonian Monte Carlo. *International Conference on Machine Learning*, 2015.
- [24] S Karlin. *A first course in stochastic processes*. 1968.
- [25] Akihisa Ichiki and Masayuki Ohzeki. Violation of detailed balance accelerates relaxation. *Physical Review E*, 88(2):020101, August 2013.
- [26] K. Balakrishnan. *Exponential Distribution: Theory, Methods and Applications*. CRC Press, 1996.
- [27] Elizabeth L Wilmer. Markov Chains and Mixing Times David A . Levin Yuval Peres. 2009.
- [28] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959, 2012.

Appendix

A Transition rates satisfy balance condition

The continuous time balance condition states that, at the steady state distribution, there is no net change in the probability of states,

$$\left. \frac{\partial p(\zeta)}{\partial t} \right|_{p(\zeta)=\pi(\zeta)} = 0. \quad (19)$$

In order to demonstrate that we satisfy the balance condition, we evaluate $\left. \frac{\partial p(\zeta)}{\partial t} \right|_{p(\zeta)=\pi(\zeta)}$ using the transition rates from Equation 9,

$$\begin{aligned} \left. \frac{\partial p(\zeta)}{\partial t} \right|_{p(\zeta)=\pi(\zeta)} &= -\pi(\zeta) \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \\ &\quad + \pi(\mathbf{L}^{-1}\zeta) \left[\frac{\pi(\mathbf{L}\mathbf{L}^{-1}\zeta)}{\pi(\mathbf{L}^{-1}\zeta)} \right]^{\frac{1}{2}} \\ &\quad - \pi(\zeta) \max \left[0, \left[\frac{\pi(\mathbf{L}^{-1}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \right] \\ &\quad + \pi(\mathbf{F}\zeta) \max \left[0, \left[\frac{\pi(\mathbf{L}^{-1}\mathbf{F}\zeta)}{\pi(\mathbf{F}\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}\mathbf{F}\zeta)}{\pi(\mathbf{F}\zeta)} \right]^{\frac{1}{2}} \right] \\ &\quad - \pi(\zeta) \beta \\ &\quad + \int d\mathbf{x}' d\mathbf{v}' \pi(\mathbf{x}', \mathbf{v}') \delta(\mathbf{x}' - \mathbf{x}) \pi(\mathbf{v}) \beta, \end{aligned} \quad (20)$$

where negative terms correspond to probability flow out of state ζ into other states, and positive terms correspond to probability flow from other states into state ζ . There are only a small number of terms because transitions are only allowed to/from a limited set of states. We now proceed to simplify and cancel terms,

$$\begin{aligned} \left. \frac{\partial p(\zeta)}{\partial t} \right|_{p(\zeta)=\pi(\zeta)} &= -[\pi(\mathbf{L}\zeta) \pi(\zeta)]^{\frac{1}{2}} + [\pi(\zeta) \pi(\mathbf{L}^{-1}\zeta)]^{\frac{1}{2}} \\ &\quad - \pi(\zeta) \max \left[0, \left[\frac{\pi(\mathbf{L}^{-1}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \right] + \pi(\zeta) \max \left[0, \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} - \left[\frac{\pi(\mathbf{L}^{-1}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \right] \\ &\quad - \pi(\zeta) \beta + \pi(\mathbf{x}) \pi(\mathbf{v}) \beta \end{aligned} \quad (21)$$

$$\begin{aligned} &= -[\pi(\mathbf{L}\zeta) \pi(\zeta)]^{\frac{1}{2}} + [\pi(\zeta) \pi(\mathbf{L}^{-1}\zeta)]^{\frac{1}{2}} \\ &\quad + \pi(\zeta) \left[\frac{\pi(\mathbf{L}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} - \pi(\zeta) \left[\frac{\pi(\mathbf{L}^{-1}\zeta)}{\pi(\zeta)} \right]^{\frac{1}{2}} \end{aligned} \quad (22)$$

$$= -[\pi(\mathbf{L}\zeta) \pi(\zeta)]^{\frac{1}{2}} + [\pi(\zeta) \pi(\mathbf{L}^{-1}\zeta)]^{\frac{1}{2}} + [\pi(\mathbf{L}\zeta) \pi(\zeta)]^{\frac{1}{2}} - [\pi(\zeta) \pi(\mathbf{L}^{-1}\zeta)]^{\frac{1}{2}} \quad (23)$$

$$= 0. \quad (24)$$

Therefore the transition rates in MJHMC satisfy the balance condition for $\pi(\zeta)$, as claimed.

B Hyperparameter search

B.1 Demonstration of optimized hyperparameters

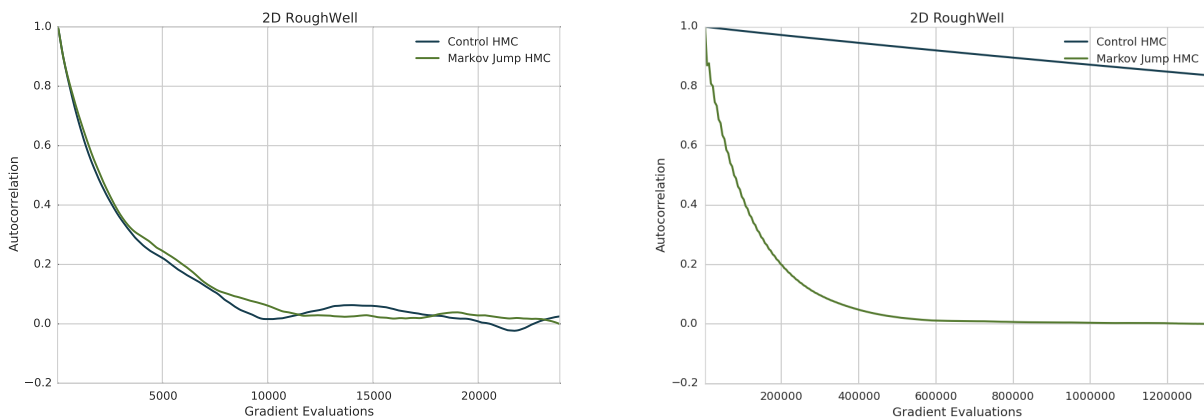


Figure A.1: Comparison of mixing performance of MJHMC and standard discrete-time HMC with both samplers set to the same hyperparameters (a) Both samplers set to $\varepsilon = 0.591686$, $\beta = 0.429956$, $M = 25$, the best setting for standard HMC found by Spearmint (b) Both samplers set to $\varepsilon = 3.0$, $\beta = 0.012314$, $M = 25$, the best settings for MJHMC found by Spearmint

The autocorrelation data illustrated in figure A.1 demonstrates that Spearmint found effective hyperparameters for MJHMC and standard discrete-time HMC on our chosen energy function. Each sampler outperforms the other when both are set its optimized setting of hyperparameters.

B.2 Illustration of Spearmint search

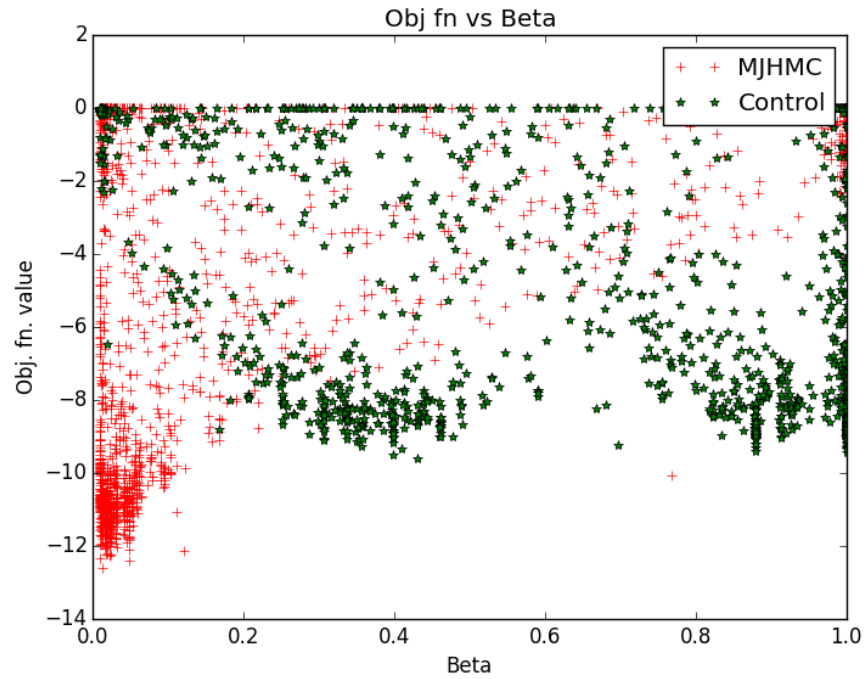


Figure A.2: Search performance projected onto the β axis.

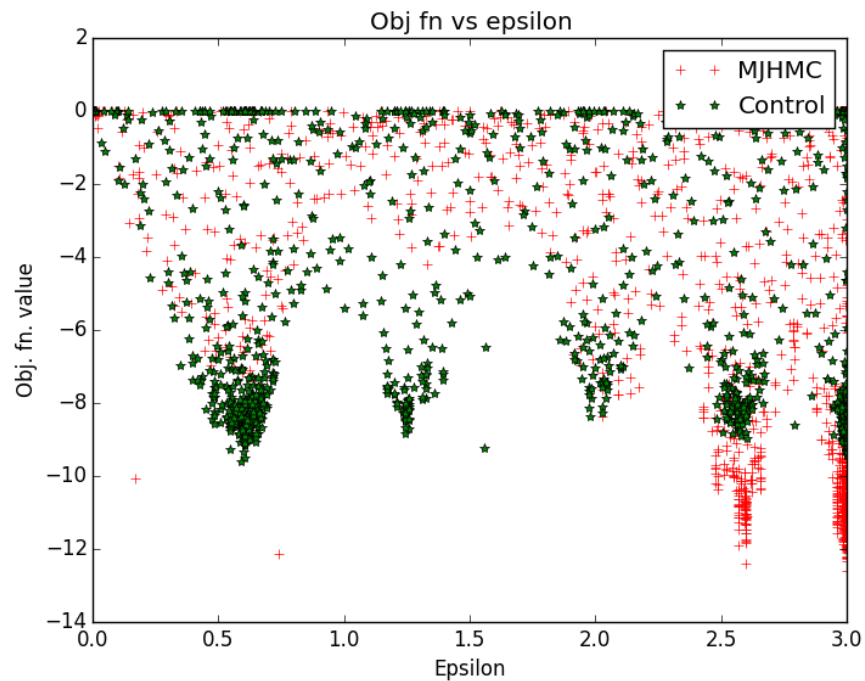


Figure A.3: Search performance projected onto the ε axis.

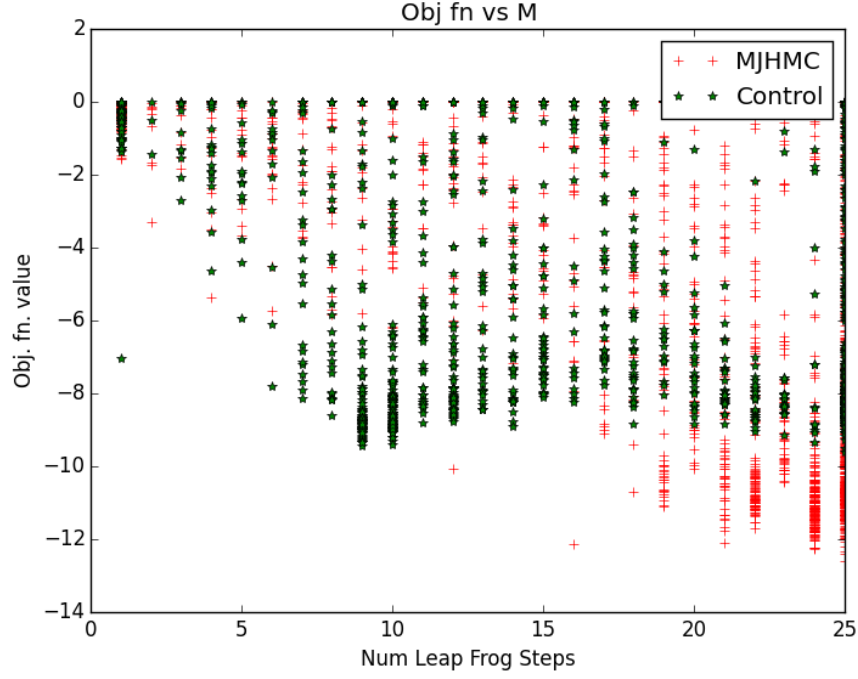


Figure A.4: Search performance projected onto the M axis.

Figures A.2, A.3, A.4 illustrate the overall structure of Spearmint’s search for hyperparameters. The green stars represent a trial setting of MJHMC hyperparameter and the red crosses represent a trial setting of standard HMC hyperparameters. The y-axis represents the value of the objective function for each trial setting. It can be seen that in A.2 that MJHMC chooses smaller β values which suggests wanting to corrupt momentum more slowly as compared to the control case. It can also be seen from ?? that MJHMC prefers larger steplengths for the integrator (ϵ) and steps (M) .

C Derivation of equation 11

First we calculate $P(\tau_2 \leq \tau_1)$ where τ_i is drawn from $\text{Exp}(\lambda_i)$ for $i = 1, 2$:

$$\begin{aligned}
P(\tau_2 \leq \tau_1) &= \int_0^\infty P(\tau_1 = t) P(\tau_2 \leq t) dt \\
&= \int_0^\infty (\lambda_1 \exp(-\lambda_1 t)) \int_0^t (\lambda_2 \exp(-\lambda_2 \tau)) d\tau dt \\
&= \lambda_1 \int_0^\infty \exp(-\lambda_1 t) [1 - \exp(-\lambda_2 t)] dt \\
&= \lambda_1 \int_0^\infty \exp(-\lambda_1 t) - \exp(-(\lambda_1 + \lambda_2)t) dt \\
&= \lambda_1 \left[-\frac{1}{\lambda_1} \exp(-\lambda_1 t) \Big|_0^\infty \right] - \left[-\frac{1}{\lambda_1 + \lambda_2} \exp(-(\lambda_1 + \lambda_2)t) \Big|_0^\infty \right] \\
&= \lambda_1 \left[\frac{1}{\lambda_1} - \frac{1}{\lambda_1 + \lambda_2} \right] \\
&= 1 - \frac{\lambda_1}{\lambda_1 + \lambda_2} \\
&= \frac{\lambda_2}{\lambda_1 + \lambda_2}
\end{aligned}$$

Let $\tau_{i,j}$ be drawn from $\text{Exp}(\Gamma_{ij})$. Then

$$\begin{aligned}
P(\zeta_j \mid \zeta_i) &= P(\tau_{i,j} = \min\{\tau_{i,1}, \tau_{i,2}, \dots, \tau_{i,n}\}) \\
&= \prod_{k \neq j}^n P(\tau_{i,j} \leq \tau_{i,k}) \\
&= \prod_{k \neq j}^n \frac{\Gamma_{i,j}}{\Gamma_{i,j} + \Gamma_{i,k}}
\end{aligned}$$